

# A Geometric Framework For Density Modeling

Sutanoy Dasgupta<sup>\*1</sup>, Debdeep Pati<sup>1</sup>, and Anuj Srivastava<sup>1</sup>

<sup>1</sup>Department of Statistics, Florida State University

## Abstract

We introduce a geometric approach for estimating a probability density function (*pdf*) given its samples. The procedure involves obtaining an initial estimate of the *pdf* and then transforming it via a warping function to reach the final estimate. The initial estimate is intended to be computationally fast, albeit suboptimal, but its warping creates a larger, flexible class of density functions, resulting in substantially improved estimation. The warping is accomplished by mapping diffeomorphic functions to the tangent space of a Hilbert sphere, a vector space whose elements can be expressed using an orthogonal basis. Using a truncated basis expansion, we estimate the optimal warping and, thus, the optimal density estimate. This framework is introduced for univariate, unconditional *pdf* estimation and then extended to conditional *pdf* estimation. The approach avoids many of the computational pitfalls associated with current methods without losing on estimation performance. In presence of irrelevant predictors, the approach achieves both statistical and computational efficiency compared to classical approaches for conditional density estimation. We derive asymptotic convergence rates of the density estimator and demonstrate this approach using synthetic datasets, and a case study to understand association of a toxic metabolite on preterm birth.

KEYWORDS: conditional density; density estimation; Hilbert sphere; irrelevant predictors; sieve estimation; tangent space; weighted likelihood maximization

---

<sup>\*</sup>Corresponding Author: s.dasgupta@stat.fsu.edu

## 1. INTRODUCTION

Estimating a probability density function (*pdf*) is an important and well studied field of research in statistics. The most basic problem in this area is that of univariate *pdf* estimation from *iid* samples, henceforth referred to as unconditional density estimation. Another problem of utmost importance is conditional density estimation. Here one needs to characterize the behavior of the response variable for each of the different values of the predictors. A natural extension of these problems is that of estimating the joint *pdf* for multiple random variables.

Given the importance of *pdf* estimation in statistics and related disciplines, a large number of solutions have been proposed for each of these problems. While the earliest works focused on parametric solutions, the trend over the last three decades is to use a nonparametric approach as it minimizes making assumptions about the underlying density (and the relationships between variables for conditional and joint densities). The most common nonparametric techniques are kernel based; refer to Rosenblatt et al. [1956], Hall et al. [1991], Sheather and Jones [1991], Li and Racine [2007] for a narrative of works. Related to these approaches are “tilting” or “data sharpening” techniques for unconditional density estimation, for example Hjort and Glad [1995], Doosti and Hall [2016] and the references therein. Kernel methods are very powerful in univariate setting. However, as the number of variables involved gets higher, the methods become computationally inefficient because of the complexities involved in bandwidth selection, especially in conditional density estimation setup. Another common approach [Leonard, 1978, Lenk, 1988, 1991, Tokdar et al., 2010, Tokdar, 2012] for *pdf* estimation is a two step estimation procedure, where one first estimates an initial *pdf*, say  $f_p$ , from the data, perhaps restricting to a parametric family. Then, one *improves* upon this estimate by considering a function  $g(x)$  depending on the initial estimate  $f_p$  and an unknown function  $W$ , to reach a final estimate of the type  $\exp\{g(x)\} / \int_x \exp\{g(x)\} dx$ . In a classical paradigm one can choose  $\exp\{g(x)\}$  as  $W(x)f_p(x)$  where  $W(x)$  is a positive function. In a Bayesian context, the function  $g$  is assigned a Gaussian process prior  $W$  with mean  $f_p$ . Often the calculation of the normalization constant makes the computation very cumbersome. On the other hand, two step procedures for conditional density estimation are based on estimating the conditional mean function, and then the conditional density of the residuals [Hansen, 2004]. Over the recent years, Bayesian methods for estimating *pdfs* based on mixture models and latent variables have received a lot of attention because of their excellent practical performances and an increasingly rich set of algorithmic tools for posterior computation using Markov Chain Monte Carlo

(MCMC) methods. References include Escobar and West [1995], Müller et al. [1996], MacEachern and Müller [1998], Kalli et al. [2011], Jain and Neal [2012], Kundu and Dunson [2014], Bhattacharya et al. [2010] among others. However these results also come at a very high computational cost associated with implementing the MCMC algorithms. Applications of flexible Bayesian models for conditional densities are discussed in MacEachern [1999], De Iorio et al. [2004], Griffin and Steel [2006], Dunson et al. [2007], Chung and Dunson [2009], Norets and Pelenis [2012], among others. Although the literature suggests that such methods based on mixture models have several attractive properties, they lack interpretability and the MCMC solutions for model fitting are overly complicated and expensive.

In this article a geometrically intuitive framework for a two step approach is proposed that is applicable to both conditional and unconditional density estimation. The motivation here is to avoid the computational cost associated with traditional techniques, while not losing out on practical performance. Assume that we have a strictly positive univariate density  $f_p$  on the interval  $[0, 1]$ ;  $f_p$  serves as an initial estimate as in the traditional two step methods. Let  $\Gamma$  be the set of all positive diffeomorphisms from  $[0, 1]$  to itself:

$$\Gamma = \{\gamma | \gamma \text{ is smooth, } \gamma^{-1} \text{ is smooth, } \dot{\gamma} > 0, \gamma(0) = 0, \gamma(1) = 1\} \quad (1)$$

The elements of  $\Gamma$  play the role of warping functions, or transformations of  $f_p$ . Given a  $\gamma \in \Gamma$ , the transformation of  $f_p$  is defined by:

$$(f_p, \gamma) = (f_p \circ \gamma) \dot{\gamma}. \quad (2)$$

Henceforth this transformation is referred to as *warping* and the resulting pdf  $f$  as a *warped density*. This mapping is comprehensive in the sense that it forms a transitive group action on the space of all strictly positive pdfs on  $[0, 1]$ . This allows one to go from any positive pdf to any other positive pdf using an appropriate  $\gamma$ . Note that since  $\int_0^1 f_p(\gamma(x)) \dot{\gamma}(x) dx = 1$ , there is no need to normalize this transformation. However, since  $\Gamma$  is not a vector space, one has to map the elements of  $\Gamma$ , using a nonlinear mapping, to a vector space to facilitate expansion using an orthogonal basis. The diffeomorphisms as transformations of a probability density shape have been used in the past, albeit with a different setup and scope [Saoudi et al., 1994, 1997]. Also, the notion of transformation between pdfs has been used in the literature on *optimal transport* [Tabak and Turner, 2013, Tabak and Trigila, 2014]. The transport is achieved using an iterated composition of maps and is thus very different from estimating a diffeomorphism.

In this article, we develop the framework for estimating an unconditional, univariate pdf defined on

$[0, 1]$ . This simple setting helps explain and illustrate the main ingredients of the framework. Besides, the proposed geometric framework is naturally univariate in the sense that the transformation defined in (2) acts on univariate density shapes making it a logical starting point for developments. In this simple setup, the approach delivers excellent performance while avoiding heavy computational cost, and is comparable to standard kernel methods, even at very low sample sizes. The framework is extended to univariate densities with unknown support by scaling the observation domain to  $[0, 1]$ . A defining characteristic of this warping transformation is that it takes any initial density estimate to any final desired estimate, as long as they are both strictly positive. Hence the initial estimate can be constructed in anyway – parametric (e.g. gaussian) or nonparametric (e.g. kernel estimate), and is allowed to be a sub-optimal estimate of the true density.

The second part of the article focuses on extending the framework to estimation of conditional density  $f(y|x)$  from  $\{(y_i, x_i) : i = 1, \dots, n, y \in \mathbb{R}, x \in \mathbb{R}^d, d \geq 1\}$ . The approach is to start with a nonparametric mean regression model of the form  $y_i = m(x_i) + \epsilon_i$ ,  $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ , where  $m(\cdot)$  is estimated using a standard nonparametric estimator, to obtain an initial conditional density estimate  $f_{p,x} \equiv \mathcal{N}(\hat{m}(x), \hat{\sigma}^2)$  at the location  $x$ . Then  $f_{p,x}$  is warped using a warping function  $\gamma_x$  into a final conditional density estimate. Importantly, the set  $\Gamma$  and the warping transformation remain the same as the unconditional density estimation, and do not involve the predictors. Only the choice of the element  $\gamma_x \in \Gamma$  varies with the predictor  $x$ . This makes the estimation process computationally efficient even in the presence of many predictors. The selection of  $\gamma_x$  is based on a weighted-likelihood objective function that borrows information from the neighborhood of the location  $x$  at which the conditional density is evaluated. Since the transformation operation does not involve the predictors, the final estimate is naturally resistant to the effect of predictors that are irrelevant to the response.

The rest of this paper is organized as follows. Section 2 introduces the formal methodology and framework on a univariate unconditional density estimation setting. Section 3 presents notations and the convergence rate of the density estimator. Section 4 focuses on the practical implementation of the method and outlines the associated numerical techniques. Illustrations of univariate density estimation on simulated datasets are presented in Section 5. In section 6 the theory for conditional density estimation is developed and the properties of the proposed method are illustrated using simulated datasets. The application of conditional density estimation using the proposed framework on a real dataset is also presented. Section 7 includes a discussion on how the ideas presented in earlier sections can be used to perform bivariate uncon-




Figure 1: *Left: The true pdf  $f_t$  is estimated by transforming an initial estimate  $f_p$  by the warping function  $\gamma$ . The larger the set of allowed  $\gamma$ s, the better the estimate is. Right: Representing warping function  $\gamma$  as element of the tangent space  $T_1(\mathbb{S}_\infty^+)$ .*

ditional density estimation, the challenges involved in higher dimensions, and a simple bivariate illustration. Finally, the proofs of the convergence rate presented in Section 3 is derived in an appendix.

## 2. PROPOSED FRAMEWORK

In this section we develop a two-step framework for estimating univariate, unconditional *pdf*, and start by introducing some notation. Let  $\mathcal{F}$  be the set of all strictly positive, univariate probability density functions on  $[0, 1]$ . Let  $f_t \in \mathcal{F}$  denote the underlying true density and  $x_i \sim f_t$ ,  $i = 1, 2, \dots, n$  be independent samples from  $f_t$ . Furthermore, let  $\mathcal{F}_p$  be a pre-determined subset of  $\mathcal{F}$ , such that an optimal element (defined appropriately)  $f_p \in \mathcal{F}_p$  is relatively easy to compute. For instance, any parametric family with simple estimation procedure is a good candidate for  $f_p$ . Similarly, kernel density estimates are also good since they are computationally efficient and robust in univariate setups.

Next, we define a warping-based transformation of elements of  $\mathcal{F}_p$ , using elements of  $\Gamma$  defined in (1). Note that  $\Gamma$  is an infinite-dimensional set that has a group structure under composition as the group operation. That is, for any  $\gamma_1, \gamma_2 \in \Gamma$ , the composition  $\gamma_1 \circ \gamma_2 \in \Gamma$ . The identity element of  $\Gamma$  is given by  $\gamma_{\text{id}}(t) = t$ , and for every  $\gamma \in \Gamma$ , there is a function  $\gamma^{-1}$  such that  $\gamma \circ \gamma^{-1} = \gamma_{\text{id}}$ . For any  $f_p \in \mathcal{F}_p$  and  $\gamma \in \Gamma$ , define the mapping  $(f_p, \gamma) = (f_p \circ \gamma)\gamma$  as given in (2). The importance of this mapping comes from the following result.

**Proposition 1.** *The mapping  $\mathcal{F} \times \Gamma \rightarrow \mathcal{F}$ , specified above, forms an action of  $\Gamma$  on  $\mathcal{F}$ . Furthermore, this*

*action is transitive.*

**Proof:** We can verify the two properties in the definition of a group action: (1) For any  $\gamma_1, \gamma_2 \in \Gamma$  and  $f \in \mathcal{F}$ , we have  $((f, \gamma_1), \gamma_2) = (((f \circ \gamma_1)\dot{\gamma}_1) \circ \gamma_2)\dot{\gamma}_2 = (f, \gamma_1 \circ \gamma_2)$ . (2) For any  $f \in \mathcal{F}$ ,  $(f, \gamma_{\text{id}}) = f$ . To show transitivity, we need to show that given any  $f_1, f_2 \in \mathcal{F}$ , there exists a  $\gamma \in \Gamma$ , such that  $(f_1, \gamma) = f_2$ . If  $F_1$  and  $F_2$  denote the cumulative distribution functions associated with  $f_1$  and  $f_2$ , respectively, then the desired  $\gamma$  is simply  $F_1^{-1} \circ F_2$ . Since  $f_1$  is strictly positive,  $F_1^{-1}$  is well defined and  $\gamma$  is uniquely specified. Furthermore, since  $f_2$  is strictly positive, we have  $\dot{\gamma} > 0$ .  $\square$

This result implies that together the pair  $(f_p, \gamma)$  spans the full set  $\mathcal{F}$ , if  $\gamma$  is chosen freely from  $\Gamma$ . One can reach any  $f_t \in \mathcal{F}$ , from any  $f_p \in \mathcal{F}_p$  using an appropriate element of  $\Gamma$ . However, if one uses  $\Gamma_J$ , a finite-dimensional submanifold of  $\Gamma$ , instead of the full  $\Gamma$ , we may not reach the desired  $f_t$ , but only get closer. This is depicted pictorially in the left panel of Figure 1 where the inner disk denotes the set  $\mathcal{F}_p$ . The increasing rings around  $\mathcal{F}_p$  represent the set  $\{(f_p, \gamma) | f_p \in \mathcal{F}_p, \gamma \in \Gamma_J\}$  for an increasing sequence of  $\Gamma_J$ s. As  $\Gamma_J$  approaches  $\Gamma$ , the corresponding approximation approaches  $f_t$ . More details are included in Section A.1.

## 2.1 Finite-Dimensional Representation of Warping Functions

Solving an optimization problem, say maximum-likelihood estimation, over  $\Gamma$  faces two main challenges. First,  $\Gamma$  is a nonlinear manifold, and second, it is infinite-dimensional. We handle the nonlinearity by forming a bijective map from  $\Gamma$  to a tangent space of the unit Hilbert sphere  $\mathbb{S}_\infty$  (the tangent space is a vector space), and infinite dimensionality by selecting a finite-dimensional subspace of this tangent space. Together, these two steps are equivalent to finding a family of finite-dimensional submanifolds  $\Gamma_J$  that can be *flattened* into vector spaces. This allows for a representation of  $\gamma$  using orthogonal basis. Once we have a finite-dimensional representation of  $\gamma$ , we can optimize over this representation of  $\gamma$  using the maximum-likelihood criterion.

Define a function  $q : [0, 1] \rightarrow \mathbb{R}$ ,  $q(t) = \sqrt{\dot{\gamma}(t)}$ , as the square-root slope function (SRSF) of a  $\gamma \in \Gamma$ . (For a discussion on SRSFs of general functions, please refer to Chapter 4 of Srivastava and Klassen [2016]). For any  $\gamma \in \Gamma$ , its SRSF  $q$  is an element of the interior of the positive orthant of the unit Hilbert sphere  $\mathbb{S}_\infty \subset \mathbb{L}^2$ , denoted by  $\mathbb{S}_\infty^+$ . This is because  $\|q\|^2 = \int_0^1 q(t)^2 dt = \int_0^1 \dot{\gamma}(t) dt = \gamma(1) - \gamma(0) = 1$ . We have a positive orthant, boundaries excluded, because by definition,  $q$  is a strictly positive function. The

mapping between  $\Gamma$  and  $\mathbb{S}_\infty^+$  is a bijection, with its inverse given by  $\gamma(t) = \int_0^t q(s)^2 ds$ . The unit Hilbert sphere is a smooth manifold with known geometry under the  $\mathbb{L}^2$  Riemannian metric [Lang, 2012]. It is not a vector space but a manifold with a constant curvature, and can be easily flattened into a vector space locally. The chosen vector space is a tangent space of  $\mathbb{S}_\infty^+$ . A natural choice for reference, to select the tangent space, is the point  $\mathbf{1} \in \mathbb{S}_\infty^+$ , a constant function with value 1, which is the SRSF corresponding to  $\gamma = \gamma_{\text{id}}(t) = t$ . The tangent space of  $\mathbb{S}_\infty^+$  at  $\mathbf{1}$  is an infinite-dimensional vector space given by:  $T_1(\mathbb{S}_\infty^+) = \{v \in \mathbb{L}^2([0, 1], \mathbb{R}) \mid \int_0^1 v(t) dt = \langle v, \mathbf{1} \rangle = 0\}$ .

Next, we define a mapping that takes an arbitrary element of  $\mathbb{S}_\infty^+$  to this tangent space. For this *retraction*, we will use the inverse exponential map that takes  $q \in \mathbb{S}_\infty^+$  to  $T_1(\mathbb{S}_\infty^+)$  according to:

$$\exp_1^{-1}(q) : \mathbb{S}_\infty^+ \rightarrow T_1(\mathbb{S}_\infty^+), \quad v = \exp_1^{-1}(q) = \frac{\theta}{\sin(\theta)}(q - \mathbf{1} \cos(\theta)), \quad (3)$$

where  $\theta = \cos^{-1}(\langle \mathbf{1}, q \rangle)$  is the arc-length from  $q$  to  $\mathbf{1}$ . The right panel of Fig. 1 shows a pictorial illustration of the mapping from  $\mathbb{S}_\infty^+$  to the tangent space  $T_1(\mathbb{S}_\infty^+)$ .

We impose a natural Hilbert structure on  $T_1(\mathbb{S}_\infty^+)$  using the standard inner product:  $\langle v_1, v_2 \rangle = \int_0^1 v_1(t)v_2(t)dt$ . It is easy to check that since  $q \in \mathbb{S}_\infty^+$ ,  $\theta = \cos^{-1}(\langle \mathbf{1}, q \rangle) < \pi/4$ , and hence  $\|v\| = \sqrt{\int_0^1 v(t)^2 dt} = \theta < \pi/4$ , where  $v = \exp_1^{-1}(q)$ . Thus the range of the inverse exponential map is not the entire  $T_1(\mathbb{S}_\infty^+)$ , but an open subset  $T_1^0(\mathbb{S}_\infty^+) = \{v \in T_1(\mathbb{S}_\infty^+) : \|v\| < \pi/4\}$ . Further, we can select any orthogonal basis  $\mathcal{B} = \{b_j, j = 1, 2, \dots\}$  of the set  $T_1(\mathbb{S}_\infty^+)$  to express its elements  $v$  by their corresponding coefficients; that is,  $v(t) = \sum_{j=1}^\infty c_j b_j(t)$ , where  $c_j = \langle v, b_j \rangle$ . The only restriction on the basis elements  $b_j$ 's is that they must be orthogonal to  $\mathbf{1}$ , that is,  $\langle b_j, \mathbf{1} \rangle = 0$ . In order to map points back from the tangent space to the Hilbert sphere, we use the exponential map, given by:

$$\exp(v) : T_1(\mathbb{S}_\infty^+) \rightarrow \mathbb{S}_\infty, \quad \exp(v) = \cos(\|v\|)\mathbf{1} + \frac{\sin(\|v\|)}{\|v\|} v. \quad (4)$$

If we restrict the domain of the exponential map to the subset  $T_1^0(\mathbb{S}_\infty^+)$ , then the range of this map is  $\mathbb{S}_\infty^+$ .

Using these two steps, we specify the finite-dimensional (approximate) representation of diffeomorphisms.

We define a composite map  $H : \Gamma \rightarrow \mathbb{R}^J$ , pictorially illustrated in Figure 2, as

$$\gamma \in \Gamma \xrightarrow{\text{SRSF}} q = \sqrt{\dot{\gamma}} \in \mathbb{S}_\infty^+ \xrightarrow{\exp_1^{-1}} v \in T_1^0(\mathbb{S}_\infty^+) \xrightarrow{\{b_j\}} \{c_j = \langle v, b_j \rangle\} \in \mathbb{R}^J. \quad (5)$$

The range of  $H$  is  $V_\pi = \{c \in \mathbb{R}^J : \|\sum_{j=1}^J c_j b_j\| < \pi/4\} \subset \mathbb{R}^J$ . Now, we define  $G : \mathbb{R}^J \rightarrow \Gamma$ , as

$$\{c_j\} \in \mathbb{R}^J \xrightarrow{\{b_j\}} v = \sum_{j=1}^J c_j b_j \in T_1(\mathbb{S}_\infty^+) \xrightarrow{\exp_1} q = \exp_1(v) \in \mathbb{S}_\infty \rightarrow \gamma(t) = \int_0^t q(s)^2 ds. \quad (6)$$




Figure 2: A graphic representation of Eqn. 5 leading to a bijective map between  $\Gamma$  and  $V_\pi$ .

If we restrict the domain of  $G$  to  $V_\pi$ , then  $G$  is invertible and its inverse is  $H$ . Since our attention can be restricted only to the set  $V_\pi$ , rather than the entire space  $\mathbb{R}^J$ , we identify the function  $G$  as  $H^{-1}$  henceforth, restricting its domain to be  $V_\pi$ . For any  $c \in V_\pi$ , let  $\gamma_c$  denote the diffeomorphism  $H^{-1}(c)$ . For any fixed  $J$ , the set  $H^{-1}(V_\pi)$  is a finite-dimensional submanifold of  $\Gamma$ , on which we pose the estimation problem. As  $J$  goes to infinity, this submanifold converges to the full group  $\Gamma$ .

With this setting, we can rewrite the estimate of the unknown density  $f_t$ , given an initial density estimate  $f_p$ , as  $\hat{f}(t) = f_p(\gamma_{\hat{c}}(t))\dot{\gamma}_{\hat{c}}(t)$ ,  $t \in [0, 1]$ , where  $\gamma_{\hat{c}} = H^{-1}(\hat{c})$  and

$$\hat{c} = \operatorname{argmax}_{c \in \mathbb{R}^J} \left( \sum_{i=1}^n \left[ \log (f_p (\gamma_c(x_i)) \dot{\gamma}_c(x_i)) \right] \right), \quad \gamma_c = H^{-1}(c). \quad (7)$$

**Choice of Basis Functions:** One can choose from a wide range of basis elements for  $T_1(S_\infty)$ . Since the proposed analysis is on  $[0, 1]$ , one can use the Fourier basis elements (excluding 1 of course). However, other bases such as splines and Legendre polynomials can also be used. In the experimental studies, we demonstrate an example using the Meyer wavelets. Meyer wavelets have attractive properties of infinite differentiability and support over all reals. Vermehren and de Oliveira [2015] provides a closed-form expression for Meyer wavelets and scale function in the time domain, which enables us to easily use the basis set for representation. However, Meyer wavelets are not naturally orthogonal to 1 and so they need to be orthogonalized first.



We have used the geometry of  $\Gamma$  to develop a natural, local flattening of this nonlinear manifold, and then to approximate it with a finite subspace. Other, seemingly simpler, choices are also possible. For instance, since any  $\gamma$  can also be viewed as a nonnegative function in  $\mathbb{L}^2$  with appropriate constraints, it may be tempting to use  $\gamma \approx \sum_{j=1}^{\infty} c_j b_j(t)$  for some orthogonal basis  $\mathcal{B} = \{b_j, j = 1, 2, \dots\}$  of  $\mathbb{L}^2[0, 1]$ . This seems easier than our approach as it avoids going through the non-linear transformations. However the problem is the presence of both  $\gamma$  and  $\dot{\gamma}$  in the final estimate. A good approximation of  $\gamma$  does not imply a good approximation of  $\dot{\gamma}$ , but the reverse holds true, as argued below.

**Proof:** Let  $t \in I_{x_0}$ .  $|\gamma(t) - \gamma_{\text{app}}(t)| = |\int_0^t \dot{\gamma}(s) ds - \int_0^t \dot{\gamma}_{\text{app}}(s) ds| \leq \int_0^t |\dot{\gamma}(s) - \dot{\gamma}_{\text{app}}(s)| ds \leq \|\dot{\gamma} - \dot{\gamma}_{\text{app}}\|_{\infty} \cdot t \leq \|\dot{\gamma} - \dot{\gamma}_{\text{app}}\|_{\infty} \cdot x_0 \leq \|\dot{\gamma} - \dot{\gamma}_{\text{app}}\|_{\infty} \square$

### 2.3 Estimation of densities with unknown support

9

### 3. CONVERGENCE RATE BOUNDS

We have represented an arbitrary *pdf* as a function of the coefficients *w.r.t* a basis set of the tangent space. We note that in order to represent the entire space  $\mathcal{F}$ , we need a Hilbert basis with infinitely many elements. However, in practice, we use only a finite number  $J$  of basis elements. Hence, we are actually optimizing over a subset of the space of density functions based on only a few basis elements and using it to approximate the true density. This subset is called the *approximating space*. Since we are performing maximum likelihood estimation over an approximating space for *pdfs*, our estimation is akin to the sieve MLE, discussed in Wong and Shen [1995].

First, we introduce some notations. Recall that  $\mathcal{F}$  is the space of all univariate, strictly positive *pdfs* on  $[0, 1]$  and zero elsewhere. Let  $\mathcal{F}_n$  be the approximating space of  $\mathcal{F}$  when using  $k_n$  basis elements for the tangent space  $T_1(\mathbb{S}_\infty^+)$ , where  $k_n$  is some function of the number of observations  $n$ . Let  $f_p \in \mathcal{F}_p \subset \mathcal{F}$  be an initial estimate, and let  $\mathcal{F}_n = \{f_p(\gamma)\dot{\gamma}, \gamma = H^{-1}(c) \mid c \in V_\pi \subset \mathbb{R}^{k_n}\}$ , where  $H$  and  $V_\pi$  are defined in Section 2.1. As  $n \rightarrow \infty, k_n \rightarrow \infty$ . So  $\mathcal{F}_n \rightarrow \mathcal{F}$  as  $n \rightarrow \infty$ . Let  $\eta_n$  be a sequence of positive numbers converging to 0. Let  $\mathcal{Y}^{(n)}$  be the space of  $n$  observed points. We call an estimator  $\hat{p} : \mathcal{Y}^{(n)} \rightarrow \mathcal{F}_n$  an  $\eta_n$  sieve MLE if

$$\frac{1}{n} \sum_{i=1}^n \log \hat{p}(Y_i) \geq \sup_{p \in \mathcal{F}_n} \frac{1}{n} \sum_{i=1}^n \log p(Y_i) - \eta_n$$

In the proposed method, the estimated *pdf* is exactly  $\sup_{p \in \mathcal{F}_n} \frac{1}{n} \sum_{i=1}^n \log p(Y_i)$ . Therefore, this estimate is a sieve MLE with  $\eta_n \equiv 0$ . Let  $p_0$  denote the true density which is assumed to be  $\beta$ -smooth for  $\beta > 0$ . The following theorem provides convergence rates of the sieve estimators.

**Theorem 1.** *Under the assumptions listed above*

$$P(\|\hat{p}^{1/2} - p_0^{1/2}\|_2 \geq \epsilon_n^*) \leq c_1 \exp\{-c_2 n(\epsilon_n^*)^2\}$$

for some constants  $c_1, c_2 > 0$ , where  $\epsilon_n^* = Mn^{-\beta/(2\beta+1)}\sqrt{\log n}$  for a large constant  $M > 0$ .

The proof of Theorem 1 is deferred to the Appendix.

### 4. IMPLEMENTATION ISSUES

In this section we outline the estimation procedure and discuss some implementation issues. We discretize density functions using a dense uniform partition,  $T = 100$  equidistant points over the interval  $[0, 1]$ . For

approximating derivatives of a function, for example  $\dot{\gamma}$  for a warping function  $\gamma$ , we use the first-order differences. The integrals are approximated using the trapezoidal method.

For optimization of the log-likelihood function according to Eqn. 7, we use the function `fminsearch` in MATLAB in our experiments. The `fminsearch` function uses a very efficient grid search technique to find the optimal values of  $\{c_j\}$  corresponding to the basis elements to approximate the optimal warping function  $\gamma$ . However, `fminsearch` function can get stuck in locally-optimal solutions in some situations. To alleviate this problem we use an iterative, multi-resolution approach as follows. We start the optimization using a small value of  $J$  with  $c = \mathbf{0}$ , the point that maps to  $\gamma_{id} \in \Gamma$  under  $H^{-1}$ . This implies a low-resolution of the search space  $\mathbb{R}^J$  and using `fminsearch` we optimize in the search space. Then, at each successive iteration we increase the resolution by increasing  $J$  and use the optimal value of  $c$  from the previous iteration as the initial condition (with the additional components set to zero) for `fminsearch`. This slow increase in  $J$  while continually improving the optimal point  $c$  performs much better in practice than using a large value of  $J$  directly in `fminsearch`.

Another important numerical issue is the final choice of  $J$ . For a fixed sample of size  $n$ , a large value of  $J$  may lead to overfitting and  $\hat{f}$  being rough. Also, a large value of  $J$  makes it harder for the search procedure to converge to an optimal solution. Efromovich [2010] and the references there in discusses different data-driven methods to choose the number of basis elements, by considering the number of basis elements itself as a parameter. We take a different data driven technique to select the desired number of basis elements. Using a predetermined maximum number of basis points, we navigate through increasing number of basis elements and at each step, we compute the value of the Akaike's Information Criterion (AIC) and choose the number of basis elements that results in the best value of the AIC, penalizing the number of basis functions used. We summarize the full procedure in **Algorithm 1**.

---

**Algorithm 1** Improving solutions using `fminsearch` by tweaking the starting points

---

- i. Start with a low number of basis elements, say  $J$
  - ii. Use  $\mathbf{0}$  vector as the starting point and find the solution  $\mathbf{d}$  using `fminsearch`.
  - iii. Increase the number of basis elements, say  $J_1$  more basis elements.
  - iv. Use  $[\mathbf{0}, \mathbf{0}]$  and  $[\mathbf{d}, \mathbf{0}]$  as two starting points. Compare the AIC for the two cases and choose the solution with better AIC value. Call the solution  $\mathbf{d}$  the optimal solution.
  - v. If the number of basis elements exceeds a predetermined large number, stop. Else go to step iii.
-

## 5. SIMULATION STUDY

Next, we present results from experiments on univariate unconditional density estimation procedure involving two simulated datasets. These computations are performed on an Intel(R) Core(TM) i7-3610QM CPU processor laptop, and the computational times are reported for each experiment. We compare the proposed solution with two standard techniques: (1) kernel density estimates with bandwidth selected by unbiased cross validation method, henceforth referred to as *kernel(ucv)*, (2) a standard Bayesian technique using the function *DPdensity* in the R package *DPPackage*. The Bayesian approach naturally has a longer run-time. For both the simulated examples, we use 2000 MCMC runs with 500 iterations as burn in period for the Bayesian technique. We compare the methods both in terms of numerical performance and computational cost. The simulation study can be divided into two parts.

1. Firstly, we illustrate the performance of the various methods using a representative simulation. We highlight the performance improvement over an (misspecified) initial parametric and nonparametric density estimate brought about by warping. For the initial parametric estimate we have chosen a normal density truncated to  $[0, 1]$  with mean and standard deviation estimated from the sample. For the initial nonparametric estimate, we used inbuilt MATLAB function *ksdensity*, referred to as *ksdensity* estimate henceforth.
2. In the second part, we focus on the average performance of the different techniques over 100 independent samples from the true density. In this part we use *ksdensity* as the initial estimate. We consider sample sizes of 25, 100 and 1000, to study the effect of  $n$  on estimation performance and computational cost. The performance is evaluated using multiple norms:  $L_2$ ,  $L_1$  norm and  $L_\infty$  norm, averaged over the 100 samples.

### 5.1 Example 1

We borrow the first example from Tokdar [2012] and Lenk [1991], where  $f_t \propto 0.75\exp(\text{rate} = 3) + 0.25\mathcal{N}(0.75, 2^2)$ , a mixture of exponential and normal density truncated to the interval  $[0, 1]$ : We generate  $n = 100$  observations to study estimation performance. Here we use Meyer wavelets as the basis set for the tangent space representation of  $\gamma$ s. We use an ad hoc choice of  $J = 15$  basis elements to approximate the tangent space. Also, we use an unpenalized log likelihood for optimization.




Figure 3: The left panel compares the warped estimate  $\hat{f}$  with other estimates when  $f_p$  is parametric. The middle panel shows the corresponding evolution of the negative of log-likelihood function during optimization. The right figure compares the warped estimate with others when  $f_p$  is *ksdensity*.

Figure 3 (left) panel shows a substantial improvement in the final warped estimate over the initial parametric estimate. Incidentally, it also does a better job in capturing the left peak as compared to the *kernel(ucv)* method. The rightmost panel displays the warped result when using *ksdensity* output as the initial estimate. It also provides solutions obtained using *kernel(ucv)* and *DPdensity*. Once again, this warped estimate provides a substantial improvement over the initial solution.

Table 1 summarizes estimation performance and computation cost for these methods at different sample sizes. It is observed that when  $n = 25$ , *kernel(ucv)* method outperforms the other two methods. However, for higher sample sizes, the warping-based method has a better overall performance. The computational cost of the proposed method, while higher than *kernel(ucv)*, is much less than the *DPdensity* for higher sample sizes. In this example, we also studied performance using the Fourier basis and the results were very similar and hence not reported here.

## 5.2 Example 2

For the second example we take Example 10 from Marron and Wand [1992], which uses a claw density:  $f_t = \frac{1}{2}\mathcal{N}(0, 1) + \sum_{l=0}^4 \frac{1}{10}\mathcal{N}(\frac{l}{2} - 1, (0.1)^2)$ . We estimate the domain boundaries as discussed in Section 2.3. Unlike the previous example, instead of fixing  $J$ , the number of tangent basis elements, we allow the algorithm to find the optimal  $J$ , with a maximum value of 40. Consequently, as can be seen in Table 2, the computation cost goes up. Additionally, we note that the cost is highest for  $n = 25$  and actually decreases as  $n$  increases. This is because for small  $n$  there is less information and it takes more time for the objective function to converge.

Method		DPDensity			Kernel(ucv)			Warping Estimate		
Size ( $n$ )	Norm	Mean	std. dev.	Time	Mean	std. dev.	Time	Mean	std. dev.	Time
25	$L_1$	37.26	8.63	4 sec	33.51	11.97	< 1 sec	39.53	9.8	5 sec
	$L_2$	5.05	0.9		4.5	1.44		4.96	1.27	
	$L_\infty$	1.64	0.21		1.44	0.47		1.34	0.53	
100	$L_1$	22.87	5.32	18 sec	21.9	5.54	< 1 sec	22.46	4.95	5 sec
	$L_2$	3.47	0.58		3.14	0.57		2.93	0.61	
	$L_\infty$	1.49	0.2		1.23	0.24		0.88	0.34	
1000	$L_1$	10.79	2.05	225 sec	11.57	2.14	< 1 sec	10.05	1.36	5 sec
	$L_2$	1.83	0.24		1.67	0.23		1.31	0.16	
	$L_\infty$	1.18	0.2		0.88	0.22		0.5	0.17	

Table 1: A comparison of the performances for mixture of exponential and normal example. The values of mean and std. deviation have been scaled by 100 for convenience.

Method		DPDensity			Kernel(ucv)			Warped		
Size	Norm	Mean	std. dev.	Time	Mean	std. dev.	Time	Mean	std. dev.	Time
25	$L_1$	39.15	6.29	4 sec	17.06	2.33	1 sec	18.28	3.3	105 sec
	$L_2$	5.46	0.48		2.09	0.3		2.41	0.43	
	$L_\infty$	1.2	0.05		0.5	0.14		0.64	0.17	
100	$L_1$	28.39	4.55	26 sec	8.54	2.38	1 sec	9.06	2.6	85 sec
	$L_2$	4.31	0.46		1.18	0.28		1.3	0.35	
	$L_\infty$	1.08	0.09		0.34	0.08		0.42	0.13	
1000	$L_1$	19.28	1.63	331 sec	2.4	0.38	1 sec	2.46	0.43	71 sec
	$L_2$	3.16	0.15		0.38	0.06		0.4	0.08	
	$L_\infty$	0.83	0.04		0.14	0.03		0.15	0.04	

Table 2: Comparison for claw density example. The values of mean and std. deviation have been scaled by 100 for convenience. The computation time for warped estimate is much higher than the previous example because **Algorithm 1** was employed to obtain the optimal number of basis elements




Figure 4: The left panel shows the improvement over initial *ksdensity* estimate. Both *kernel(ucv)* and *warped* estimate have a good performance here. The right panel shows that all the methods fail to capture all the peaks. *Kernel(ucv)* performance is very similar to the *warped* estimate.

Table 2 shows that at  $n = 1000$ , the performances of all three methods are similar, especially between *kernel(ucv)* and *warped* density estimate. In fact, the *warped* density estimate and *kernel(ucv)* perform similarly even at low sample sizes, while *DPdensity* performs poorly. These results were obtained using the Fourier basis but the results for Meyer basis were similar.

## 6. EXTENSION TO CONDITIONAL DENSITY ESTIMATION

The idea of using diffeomorphisms to warp an initial density estimate, while maximizing likelihood, extends naturally to conditional density estimation. Consider the following setup: Let  $X$  be a fixed  $d$ -dimensional random variable with a positive density on its support. Let  $Y \sim f_t(m(X), \sigma_X^2)$ , where  $f_t$  is the unknown conditional density that changes smoothly with  $X$ ;  $m(X)$  is the unknown mean function, assumed to be smooth; and,  $\sigma_X^2$  is the unknown variance, which may or may not depend on  $X$ .  $Y$  is assumed to have a univariate, continuous distribution with support on unknown interval  $[A, B]$ . (In the simulation examples we have assumed  $X \sim U[0, 1]^d$  for simplicity.) We observe the pairs  $(Y_i, X_i), i = 1, \dots, n$ , and are interested in recovering the conditional density  $f_t(m(X), \sigma^2)$ .

In order to initialize estimation, we assume a nonparametric mean regression model of the form  $y_i = m(x_i) + \epsilon_i$ ,  $\epsilon_i \sim f_p(0, \sigma^2)$ , where  $m(\cdot)$  is estimated using standard local linear regression,  $f_p$  is an initial estimate for the conditional density of the response variable, and  $\sigma^2$  is estimated using the sample standard deviation of the residuals  $Y_i - \hat{m}(X_i)$ . We have used truncated normal density as  $f_p$  in the experiments presented later but other choices are equally valid. Let  $F_{p, x_0}$  be the corresponding initial estimate of the con-

ditional distribution function of  $Y$ , given  $X = x_0$  for some given value of the predictor  $x_0$ . Then, the warped density estimate, for a warping function  $\gamma$  and location  $x_0$ , is  $f_{w,x_0}(y|X = x_0) = f_p(\gamma(y), \hat{m}(x_0), \hat{\sigma}^2)\dot{\gamma}(y)$ . If  $F_{t,x_0}$  is the true conditional distribution function of  $Y$ , given  $X = x_0$ , then the optimal  $\gamma$  at location  $x_0$  is  $\gamma_{x_0} = F_{p,x_0}^{-1} \circ F_{t,x_0}$ .

In case the support of  $Y$  is not  $[0, 1]$ , we can simply rescale using  $Z = (Y - A)/(B - A)$  and accordingly scale the mean and variance estimate. Here  $A$  and  $B$  are the boundaries chosen as described in section 2.3. Setting  $f_{p,x_0} \equiv f_p(\hat{m}(X), \hat{\sigma}^2)$ , we estimate the optimal  $\gamma$  by a weighted maximum likelihood estimation:

$$\hat{\gamma}_{x_0} = \operatorname{argmax}_{\gamma \in \Gamma} \left( \sum_{i=1}^n \log(f_{p,x_0}(\gamma(y_i)|x_i)\dot{\gamma}) W_{x_0,i} \right),$$

where  $W_{x_0,i}$  is the localized weight associated with the  $i$ th observation, calculated according to:

$$W_{x_0,i} = \frac{\mathcal{N}(\|X_i - x_0\|_2/h(x_0); 0, 1)}{\sum_{j=1}^n \mathcal{N}(\|X_j - x_0\|_2/h(x_0); 0, 1)}$$

where  $\mathcal{N}(\cdot, 0, 1)$  is the standard normal *pdf* and  $h(x_0)$  is the parameter that controls the relative weights associated with the observations. This parameter is akin to the bandwidth parameter associated with traditional kernel methods for density estimation. A very large value of  $h(x_0)$  distributes approximately equal weight to all the observations, whereas a very small value considers only the observations in a neighborhood around  $x_0$ . Since  $h(x_0)$  is scalar, the tremendous computational cost associated with obtaining cross-validated bandwidths in each predictor dimension, when the predictor dimension is high, is avoided. When the predictor is one-dimensional, the parameter  $h(x_0)$  is chosen according to the location  $x_0$  using a two-step procedure as follows:

1. Compute a standard kernel density estimate  $\hat{K}$  of the predictor space using a fixed bandwidth chosen according to any standard criterion. For our purposes, we simply used the `ksdensity` estimate inbuilt in MATLAB which chooses the bandwidth optimal for normal densities. Let  $h$  be the fixed bandwidth used.
2. Then, set the bandwidth parameter  $h(x_0)$  at location  $x_0$  to be  $h(x_0) = h/\sqrt{\hat{K}(x_0)}$ .

The intuition is that  $h$  controls the overall smoothing of the predictor space based on the sample points, and the  $\sqrt{\hat{K}(x_0)}$  stretches or shrinks the bandwidth at the particular location. At a sparse region, increased borrowing of information from the other data points is desirable in order to reduce the variance of the estimate,



whereas in dense regions a reduced borrowing of information from far away points reduces the bias of the density estimates. A location from a sparse region is expected to have a low density estimate, and a location from a dense region is expected to have a high density estimate. Hence, varying the bandwidth parameter inversely with the density estimate helps adapt to the sparsity around the point of interest. The choice of the adaptive bandwidth parameter is motivated from the variable bandwidth kernel density estimators discussed in Terrell and Scott [1992], Van Kerm et al. [2003] and Abramson [1982], among others. In case of  $d$  independent predictors, the bandwidth parameter  $h(\mathbf{x}_0)$  at  $\mathbf{x}_0$  is chosen as follows:

1. Compute the kernel density estimate  $\hat{K}_i, i \in 1, \dots, d$  for the  $d$  predictors with associated bandwidths  $h_1, h_2, \dots, h_d$ . Then  $h$  is chosen as the harmonic mean of the  $h_i$ 's.
2. Once  $h$  is obtained, the bandwidth parameter  $h(\mathbf{x}_0)$  at  $\mathbf{x}_0$  is given by:

$$h(\mathbf{x}_0) = h / \left( \prod_{i=1}^d \sqrt{\hat{K}_i(x_{0i})} \right) \quad (8)$$

where  $x_{0i}$  is the  $i$ th coordinate of  $\mathbf{x}_0$ .

This choice of using the harmonic mean is based on the dependence of the minimax rates of convergence of estimators to the harmonic mean of the smoothness of the density along the different dimensions, as discussed in Lepski et al. [2015].

As was the case in unconditional *pdf* estimation, it is not required that the initial estimate to have a mean function close to the true mean function, or assume any particular form. The only requirement is that the initial conditional density should be continuous and bounded away from 0, and the density should vary smoothly with  $X$ . We use local-linear regression when the number of predictors is very small, but suggest using Nadaraya-Watson type estimator for the mean function when the number of predictors is large for computational efficiency. Finally, we note that since the optimization criterion itself depends on  $X$  through the associated weights, the choice of  $\gamma$  is influenced by the location of predictor. However, the subsequent transformation does not involve  $X$ . This helps the method avoid heavy computational cost even in the presence of many predictors.

## 6.1 Simulation studies

We present some examples to illustrate the proposed method and compare it with a standard R function `cde` in the package `hdrcode`. In the examples that follow, we assume that each of the predictors follow a  $U(0, 1)$

distribution for simplicity, but we reiterate that we do not need any restriction on the density of  $X$  except that it should be positive on its support.

In these experiments we have used a gaussian family for  $f_p$ , the initial parametric conditional density estimate. To estimate the mean function, we have used a local-linear regression function with gaussian kernel weights and bandwidth obtained from MATLAB function `ksdensity`. We skipped choosing the number of basis elements using **Algorithm 1** and fixed  $J = 6$  basis elements.

In using `hsrcde`, we also set the estimation procedure to be local-linear fitting. For comparison, we used  $n = 100$  samples to obtain a mean integrated squared-error loss function estimate, a mean absolute error estimate and a mean  $L_\infty$  loss function estimate. The responses were evaluated over a grid of 100 points at 11 equidistant locations over the support of each of the predictors. The results are summarized in Table 3.

**Example 1:** In this experiment, the true conditional density is given by:  $f(y_i|X = x_i) = \mathcal{N}(y_i; (2x_i - 1)^2, (0.1)^2)$  and  $X_i$ 's are generated according to  $U(0, 1)$ . The first row of Figure 5 shows that the initial density estimate

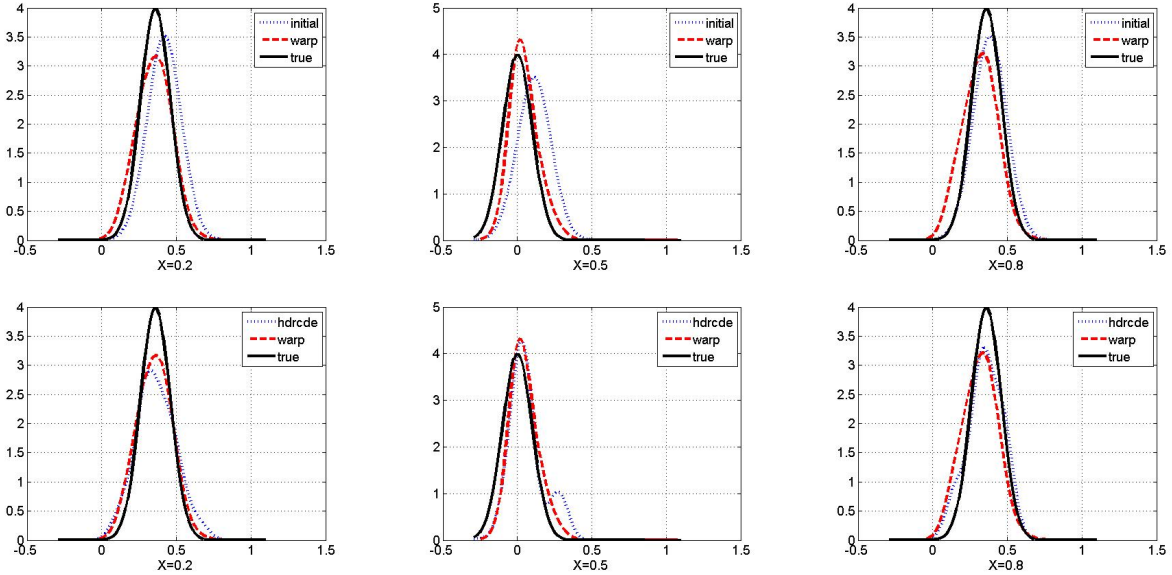


Figure 5: The first row compares the initial estimate shape with the final warped estimate at different locations of the predictor. The second row compares the performance of warped estimate with `hsrcde` at the same locations.

$f_p$  itself is very close to the truth, and the warping result is almost identical to  $f_p$ . The second row illustrates the comparison between `hsrcde` and warped estimate. The mean loss functions over 100 samples and the

results are presented in the first row of Table 3 . It suggests that the warped density estimate has better performance using all the three loss functions.

**Example 2:** As a second example, we consider a situation where the initial parametric model is gaussian family but the true conditional density is a Laplace distribution  $f(y_i|X = x_i) = \text{DExp}(y_i; \text{mean}=(2x_i - 1)^2, \text{var}=1)$




Figure 6: The first row compares the initial estimate with the final warped estimate against the true density shape at different locations of the predictor. The second row compares the performance of warped estimate with standard R technique at the same locations.

The first row of Figure 6 shows the improvement of the warped estimate over the initial gaussian shape. The standard R technique and the warped estimate has similar results illustrated by the second row of Figure 6, and summarized in the middle row of Table 3.

**Example 3:** As the third example, we consider a bivariate setup where the conditional density is a mixture model with the mixing proportion depending on one covariate and the mean functions depending on another covariate:

$$f(y_i|X = (x_{1i}, x_{2i})) = (1 - e^{-x_{2i}})\mathcal{N}(y_i; (2x_{1i} + 3), (0.5)^2) + (e^{-x_{2i}})\text{DExp}(y_i; (2x_{1i} - 1)^2, 1)$$

First row of Figure 7 shows the improvement of the warped estimate over the initial estimate at three

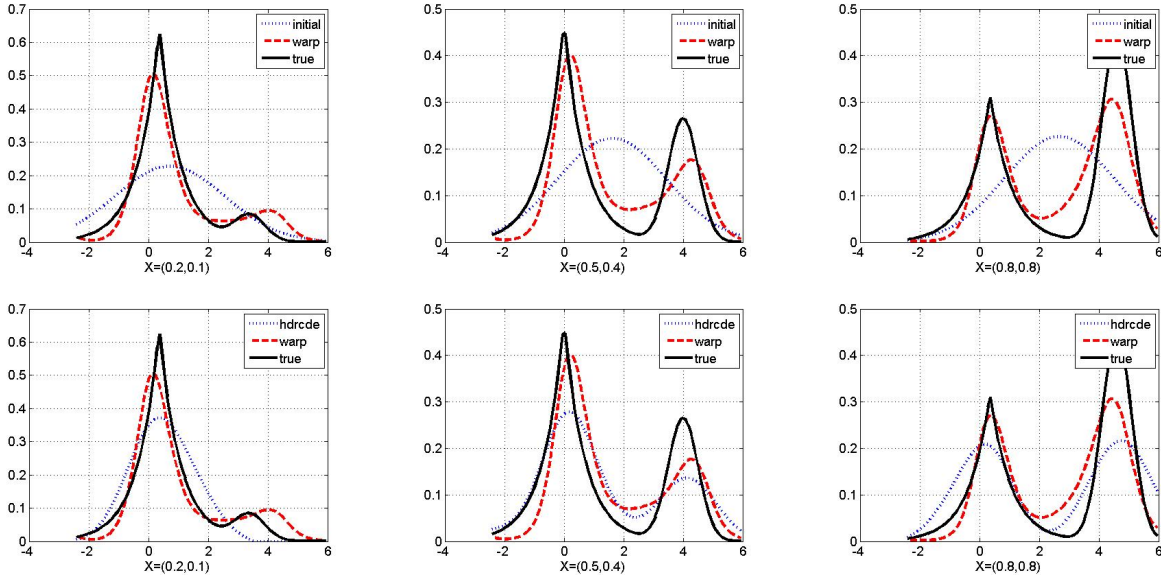


Figure 7: The first row compares the initial estimate with the final warped estimate against the true density shape at different locations of the predictor. The second row compares the performance of warped estimate with standard R technique at the same locations.

different locations on the support of  $X$ . Once again, the performance of `hdcde` and warped estimate are similar, indicated in the last row of Table 3, and in the second row of Figure 7.

Example no.	Method	MISE	Mean $L_1$ error	Mean $L_\infty$ error
Example 1	<code>hdcde</code>	0.715	0.429	0.025
	Warped	0.324	0.318	0.017
Example 2	<code>hdcde</code>	0.034	0.294	0.017
	Warped	0.034	0.287	0.014
Example 3	<code>hdcde</code>	0.057	0.491	0.021
	Warped	0.074	0.549	0.021

Table 3: A comparison of the performances of `hdcde` and warped estimate for the simulated examples.

## 6.2 Relative immunity to irrelevant predictors

The fact that given the initial density estimate shape and an appropriate choice of  $\gamma$ , the transformation defined in (2) does not involve  $X$  makes the optimization naturally resistant to the effect of irrelevant predictors, without any additional computational cost. This is very useful in situations when there are many predictors but only a small number of them are relevant to the response variable.

When the number of predictors are high, standard nonparametric kernel methods become intractable due to the huge computation cost associated with computing the bandwidth along each dimension via cross validation. Since we compute the naive `ksdensity` along each dimension, it saves the computation cost at the expense of some accuracy about the overall smoothness of the data. However, the adaptive nature of the proposed bandwidth selection helps negate some of the loss in accuracy by tweaking the bandwidth parameter according to location.

To illustrate these properties of the proposed estimator, we perform some more simulation study. We take the second and third examples presented in Section 6.1 in the same setup, and add 19 and 18 irrelevant independent predictors respectively, each distributed as  $U(0, 1)$ , resulting in 20 predictors in total. Then we fix the relevant predictors to a particular location and use 10 randomly generated locations of the irrelevant predictors on the support to study the effect of the irrelevant predictors on the resultant density estimate shape. Figure 8 illustrates the resistance of the shape to the irrelevant predictors. The density estimates are almost indistinguishable from each other in almost all cases, even though the locations of the irrelevant predictors were all different. The three panels in each row consider different locations of the relevant predictor. The entire estimation procedure takes less than 6 seconds to run, which shows the computational efficiency of the method. However, we should mention that using the **Algorithm 1** to obtain a desired number of basis elements to use will naturally increase the computation time.

### 6.3 Application to epidemiology

[Longnecker et al., 2001] studied the association of DDT metabolite DDE exposure and preterm birth in a study based on the US Collaborative Perinatal Project (CPP). DDT is very effective against malaria inflicting mosquitoes and hence is frequently used in malaria-endemic areas in spite of evidence that suggests associated health risks. Both [Longnecker et al., 2001] and [Dunson and Park, 2008] concluded that higher levels of DDE exposure is associated with higher risks of preterm birth. The response variable in question is the gestational age at delivery (GAD), and deliveries occurring prior to 37 weeks of gestation is considered as preterm. [Longnecker et al., 2001] also recorded the serum triglycerine level, among several other factors, and included it in their model since serum DDE level can be affected by concentration of serum lipids.

We study the Longnecker data to investigate the effect of varying levels of DDE on the distribution of GAD, focusing on the left tail of distribution to assess the effect on preterm births. In our study, following

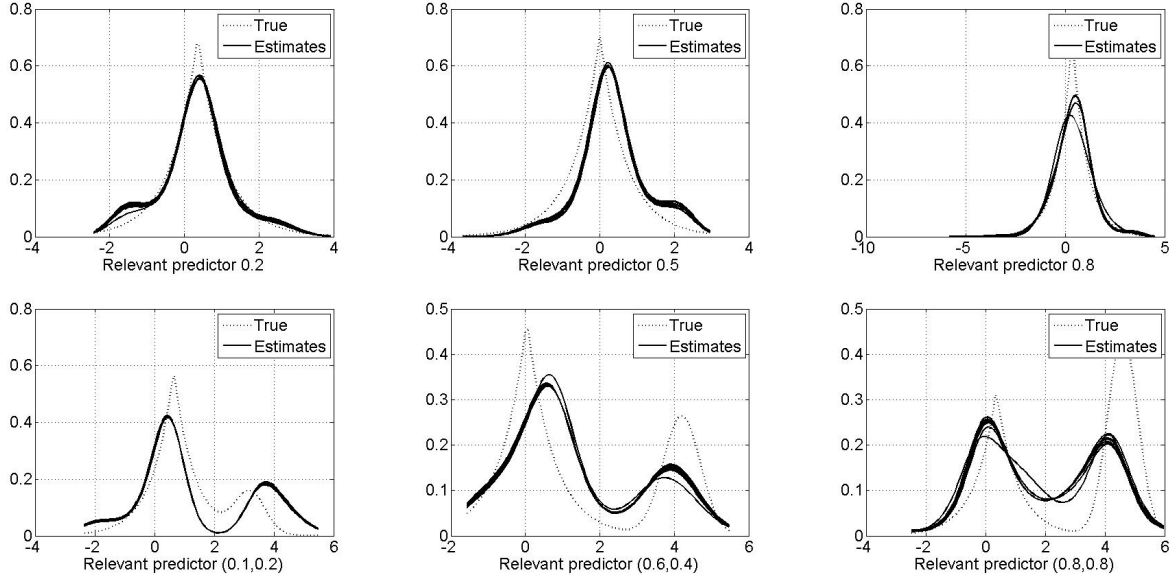


Figure 8: The figure compares the estimate shapes in the presence of irrelevant predictors at different locations of the relevant predictor(s). The solid lines correspond to different values of the irrelevant predictors. The first row is based on the example presented in Section 6.1. The second row uses the example presented in section 6.1

[Dunson and Park, 2008], we include only the 2313 subjects for whom the gestation age at delivery is less than 45 weeks, attributing higher values to measurement errors. We study the conditional density of GAD given different doses of DDE in the serum. We also study the effect of different levels of triglyceride on GAD. However, since DDE is a possible confounding factor, we conduct a bivariate analysis, including both DDE dose and triglyceride level as the covariates and study the effect on GAD at varying levels of one covariate, keeping the other fixed. We also investigate whether different levels of one covariate affect the distribution of the other.

Based on our findings, the very erratic behavior at locations where the DDE dose or triglyceride levels are 99th percentile is seen with some skepticism because of the sparsity of the data in that region. We notice an increasingly prominent peak near the left tail of GAD distribution with increasing dose of DDE, which agrees with the results of [Longnecker et al., 2001] and [Dunson and Park, 2008], shown in the left panel of Figure 9. The right panel of Figure 9 suggests a tendency of higher risks of preterm birth at higher doses of triglycerides as well, though the difference was less pronounced.

To investigate whether the results corresponding to triglycerides were confounded by the DDE doses,




Figure 9: *Distribution of gestation age at delivery for varying levels of DDE and triglyceride*

we first study the effect of triglyceride levels on DDE distribution and vice versa. Figure 10 shows that the distributions of the covariates are completely identical for varying levels of the other. The only exception is at 99th percentile of triglyceride for which the distribution of DDE doses seem to be shifted to the right. For fixed levels of triglyceride, increasing DDE doses shows an increasing left peak except where both DDE and triglyceride levels are very high, shown in Figure 11. For fixed doses of DDE the distribution of GAD at different levels of triglyceride do not follow any increasing trend and are almost indistinguishable from each other for all the different doses of DDE, as seen in Figure 12. This suggests that the increased risk of preterm birth can be attributed primarily to DDE doses, and there is no significant effect of different triglyceride levels on the gestation age. The apparent increasing risk of preterm birth for increasing level of triglycerides seen in the right panel of Figure 9 is mainly caused by DDE doses acting as a confounding factor.




Figure 10: *Distribution of DDE and triglyceride at different levels of the other*

## 7. EXTENSION TO MULTIVARIATE DENSITY ESTIMATION

The main issue with extension of the geometric approach to multivariate density estimation is that the geometry of multidimensional diffeomorphisms gets complicated. The warping function works on a univariate




Figure 11: *Distribution of gestation at varying levels of DDE for fixed values of triglyceride*




Figure 12: *Distribution of gestation at varying levels of triglyceride for fixed values of DDE*

density estimate, which forces density estimation to be restricted to univariate cases. Also for conditional density estimation, the method is restricted to univariate response variables. Interestingly, there is some existing work on density estimation in two dimensions using the same geometric idea. [Anderes and Coram, 2011, 2012] start with the group of two dimensional diffeomorphisms, but ultimately restrict to quasiconformal maps in order to develop a variational approach for density estimation. However we can avoid the very idea of two dimensional diffeomorphisms by taking a sequential approach with the tools developed on univariate density estimation and univariate response conditional density estimation to perform bivariate density estimation. For example, let  $f(X_1, X_2) = f_2(X_2|X_1)f_1(X_1)$ , where  $f$  is the joint density of  $X_1$ , and  $X_2$ ,  $f_2$  is the conditional density of  $X_2$  given  $X_1$  and  $f_1$  is the univariate density of  $X_1$ . We cannot directly estimate  $f$  using the proposed approach, but we can estimate  $f_1$  and  $f_2$  sequentially. This approach avoids the complexity of two dimensional diffeomorphisms. The only assumption needed for this approach is that the densities should be strictly positive on their support. We have illustrated bivariate density estimation with a simple bivariate example. We have generated 100 observations of independent bivariate data following mixture normal distribution:

$$X_1 \sim 0.5\mathcal{N}(-1, 0.5^2) + 0.5\mathcal{N}(1, 0.4^2), \quad X_2 \sim 0.5\mathcal{N}(1, 0.3^2) + 0.5\mathcal{N}(0, 0.5^2).$$

For this, we perform univariate unconditional density estimation of  $X_1$  and then perform conditional density estimation of  $X_2|X_1$ . The support boundaries for each variable were estimated as section 2.3. Figure 13



shows the true and estimated densities. The right panel is the estimated density which captures the peaks and the locations correctly.




Figure 13: The left panel shows the true bivariate density. The right panel shows the estimated density

Extension to higher dimensional case is not straightforward. In the case  $d = 3$ ,

$$f(X_1, X_2, X_3) = f_3(X_3|X_2, X_1)f_2(X_2|X_1)f_1(X_1),$$

the performance of  $\hat{f}_3$  deteriorates. This is because, the term  $\left(\prod_{i=1}^d \sqrt{\hat{K}_i(x_{0i})}\right)$  defined in (8) becomes an inaccurate measure of the sparsity around the point of interest  $\mathbf{x}_0$  due to the dependence structure among the predictors.

#### A. TECHNICAL DETAILS

To control the approximation error, Wong and Shen [1995] introduces a family of discrepancies. They define  $\delta_n(p_0, \mathcal{F}_n) = \inf_{q \in \mathcal{F}_n} \rho(p_0, q)$ , called the  $\rho$ -approximation error at  $p_0$ . Here  $p_0$  is the true density which is assumed to be  $\beta$ -smooth. The control of the approximation error of  $\mathcal{F}_n$  at  $p_0$  is necessary for obtaining results on the convergence rate for sieve MLEs. We follow Wong and Shen [1995] to introduce a family of indexes of discrepancy in order to formulate the condition on the approximation error of  $\mathcal{F}_n$ . Let

$$g_\alpha(x) = \begin{cases} (1/\alpha)[x^\alpha - 1], & -1 < \alpha < 0 \text{ or } 0 < \alpha \leq 1 \\ \log x, & \text{if } \alpha = 0+ \end{cases}$$

Set  $x = p/q$  and define  $\rho_\alpha(p, q) = E_p g_\alpha(X) = \int p g_\alpha(p/q)$ . We define  $\delta_n(\alpha) = \inf_{q \in \mathcal{F}_n} \rho_\alpha(p_0, q)$ . We call a finite set  $\{(f_j^L, f_j^U), j = 1, \dots, N\}$  a Hellinger  $u$ -bracketing of  $\mathcal{F}_n$  if  $\|f_j^{L1/2} - f_j^{U1/2}\|_2 \leq u$  for  $j = 1, \dots, N$ , and for any  $p \in \mathcal{F}_n$ , there is a  $j$  such that  $f_j^L \leq p \leq f_j^U$ . Let  $H(u, \mathcal{F}_n)$  be the Hellinger

metric entropy of  $\mathcal{F}_n$ , defined as the cardinality of the  $u$ -bracketing of  $\mathcal{F}_n$  of the smallest size. Let  $f_p$  be the initial estimate on which we use the group action of the space of diffeomorphisms to arrive at the final estimate. Throughout,  $c_1$  and  $c_2$  have been used to represent coefficient vectors in the tangent space of the Hilbert sphere for some fixed basis set corresponding to warping function that act on  $f_p$ . When  $c_1$  denotes the coefficient vector corresponding to the true density denoted by  $p_0 \in \mathcal{F}$  and  $c_2$  corresponds to the estimate  $q \in \mathcal{F}_n$ ,  $c_1^{>k_n}$  represents the  $(k_n + 1)$ th onwards coordinates of  $c_1$ .  $l_1, l_2, l_3$  and  $l_4$  are used to indicate specific constants. Also,  $M_1, M_2, M_3, \dots$ , have been used to represent generic constants whose value can change from step to step but is independent of other terms in the expressions.

#### A.1 Equivalence of $pdf$ space and coefficient space

Let  $f_1$  and  $f_2$  be two  $pdfs$  on  $\mathcal{F}_n$  with corresponding cumulative distribution functions  $F_1$  and  $F_2$ . Let  $f_p$  be the initial density estimate on  $\mathcal{F}_p$  such that  $f_p$  is strictly positive and Lipschitz continuous with cumulative distribution function  $F_p$ . Let  $\gamma_1 = F_p^{-1} \circ F_1$  and  $\gamma_2 = F_p^{-1} \circ F_2$ . Let  $c_1 = (c_{11}, \dots, c_{1k_n})^T$  and  $c_2 = (c_{21}, \dots, c_{2k_n})^T$  be the coefficients associated with the two elements of  $T_1(\mathbb{S}_\infty)$  corresponding to the tangent space representation of  $\gamma_1$  and  $\gamma_2$ . Here  $\mathcal{F}_n$  and  $k_n$  are as introduced in section 3. Then the following Lemma bounds the norm difference of  $f_1$  and  $f_2$  with the norm difference in the coefficients.

**Lemma 1.**  $|f_1 - f_2| \leq M_0 \|c_1 - c_2\|_1$  where  $M_0 > 0$  is a constant.

*Proof.* Let  $c_1$  and  $c_2$  be the coefficients associated with two elements of  $T_1(\mathbb{S}_\infty)$ . Then there exists  $M_1 \in \mathbb{R}$  such that  $|B_i| < M_1$ , where  $B_i$  is the  $i$ th basis function,  $i = 1, 2, \dots, k_n$ . Let  $v_1 = \sum_{i=1}^{k_n} c_{1i} B_i$ ,  $v_2 = \sum_{i=1}^{k_n} c_{2i} B_i$ . Then  $v_1, v_2 \in T_1(\mathbb{S}_\infty)$  with  $\|v_1\| < \pi/4$  and  $\|v_2\| < \pi/4$ . Hence we have

$$(v_1 - v_2)(t) = \sum_{i=1}^{k_n} (c_{1i} - c_{2i}) B_i(t) < M_1 \sum_{i=1}^{k_n} |c_{1i} - c_{2i}| < M_1 \|c_1 - c_2\|_1$$

$$\|v_1 - v_2\| = \sqrt{\int_0^1 (v_1 - v_2)^T (v_1 - v_2) dt} < M_3 \sqrt{\sum_{i=1}^{k_n} (c_{1i} - c_{2i})^2} < M_2 \|c_1 - c_2\|_2 < M_1 \|c_1 - c_2\|_1$$

Next since  $x \mapsto \|x\| = \sqrt{\int_0^1 x^2(t) dt}$  and  $x \mapsto \cos(x)$  are Lipschitz continuous, we have

$$|\cos \|v_1\| - \cos \|v_2\|| < M_2 ||v_1\| - \|v_2\|| < M_1 \|c_1 - c_2\|_1 \quad (\text{A.1})$$

Next note that  $x \mapsto \sin(x)/x$  is Lipschitz continuous. Hence we have

$$\left\| \frac{\sin \|v_1\|}{\|v_1\|} - \frac{\sin \|v_2\|}{\|v_2\|} \right\| < M_2 \|\|v_1\| - \|v_2\|\| < M_1 \|c_1 - c_2\|_1 \quad (\text{A.2})$$

Noting that

$$|q_1(t) - q_2(t)| < |\cos \|v_1\| - \cos \|v_2\|| + \left| \frac{\sin \|v_1\|}{\|v_1\|} v_1(t) - \frac{\sin \|v_2\|}{\|v_2\|} v_2(t) \right|$$

we have, combining equations A.1 and A.2,

$$\|q_1 - q_2\|_1 < M_1 \|c_1 - c_2\|_1 \quad (\text{A.3})$$

Now consider  $Q = q^2$ . Observe that

$$\begin{aligned} (Q_1 - Q_2)(t) &= q_1^2(t) - q_2^2(t) = (q_1(t) - q_2(t))(q_1(t) + q_2(t)) \\ &= (\cos \|v_1\| + \cos \|v_2\| + \frac{\sin \|v_1\|}{\|v_1\|} v_1(t) + \frac{\sin \|v_2\|}{\|v_2\|} v_2(t))(q_1(t) - q_2(t)). \end{aligned}$$

Now  $(\cos \|v_1\| + \cos \|v_2\| + \frac{\sin \|v_1\|}{\|v_1\|} v_1(t) + \frac{\sin \|v_2\|}{\|v_2\|} v_2(t))$  is a bounded function. Hence  $\|Q_1 - Q_2\|_1 < M_1 \|c_1 - c_2\|_1$  using equation A.3. Now we have  $\gamma_i(t) = \int_0^t Q_i(u) du$ ,  $t \in [0, 1]$ ,  $i = 1, 2$ . Then

$$|\gamma_1(t) - \gamma_2(t)| = \left| \int_0^t (Q_1(u) - Q_2(u)) du \right| < \int_0^t |Q_1(u) - Q_2(u)| du \leq \|Q_1 - Q_2\|_1$$

Since  $f_p$  is Lipschitz continuous and strictly positive density on  $[0, 1]$ , we have

$$\|f_p(\gamma_1) - f_p(\gamma_2)\|_1 < M_4 \|\gamma_1 - \gamma_2\|_1$$

Consider  $|f_1 - f_2| = |f_p(\gamma_1) \cdot \dot{\gamma}_1 - f_p(\gamma_2) \cdot \dot{\gamma}_2|$ . Keeping in mind that  $Q = \dot{\gamma}$ , we have

$$\begin{aligned} |f_1(t) - f_2(t)| &= |f_p(\gamma_1(t)) \cdot Q_1(t) - f_p(\gamma_2(t)) \cdot Q_2(t)| \\ &= |f_p(\gamma_1(t)) \cdot Q_1(t) - f_p(\gamma_2(t)) \cdot Q_1(t) + f_p(\gamma_2(t)) \cdot Q_1(t) - f_p(\gamma_2(t)) \cdot Q_2(t)| \\ &\leq |Q_1(t)| M_1 \|\gamma_1 - \gamma_2\|_1 + |f_p(\gamma_2(t))| \|Q_1 - Q_2\|_1 \\ &\leq M_2 \|\gamma_1 - \gamma_2\|_1 + M_3 \|\gamma_1 - \gamma_2\|_1 < M_0 \|c_1 - c_2\|_1. \end{aligned}$$

Therefore we have  $|f_1 - f_2| < M_0 \|c_1 - c_2\|_1$  for some fixed  $M_0 > 0$ . □

Note that though we worked with the  $L_1$  norm, the above relation also holds for other metrics like Hellinger metric which we use in later results.

## A.2 Probability statements on the likelihood ratio surfaces

The following theorem is based on likelihood ratio surfaces similar to Wong and Shen [1995].

**Theorem 2.** *There exists positive constants  $C_i$ ,  $i = 1, \dots, 4$  such that for any  $\epsilon > 0$ , if*

$$\int_{\epsilon^2/2^8}^{\sqrt{2}\epsilon} H^{1/2}\left(\frac{u}{C_3}, \mathcal{F}_n\right) du \leq C_4 n^{1/2} \epsilon^2 \quad (\text{A.4})$$

then

$$P^* \left( \sup_{\{\|p^{1/2} - p_0^{1/2}\|_2 \geq \epsilon, p \in \mathcal{F}_n\}} \prod_{i=1}^n p(Y_i)/p_0(Y_i) \geq \exp(-C_1 n \epsilon^2) \right) \leq 4 \exp(-C_2 n \epsilon^2)$$

*Proof.* The  $\delta$ -cover of a set  $T$  wrt a metric  $\rho$  is a set  $\{\Theta^1, \dots, \Theta^N\} \subset T$  such that for each  $\Theta \in T$ , there exists some  $i \in \{1, \dots, N\}$  with  $\rho(\Theta, \Theta_i) \leq \delta$ . The covering number  $N$  is the cardinality of the smallest delta cover. Then  $\log(N)$  is the metric entropy for  $T$ . First we bound the metric entropy for  $\mathcal{F}_n$ . Let us consider a fixed  $f_0 \in \mathcal{F}_n = f_p(\Gamma(c_0))$ . We choose the Hellinger metric so that we can borrow results directly from Wong and Shen [1995]. We note that  $H(f_1, f_2) \leq l_1 \|c_1 - c_2\|_\infty$  for some  $l_1 > 0$  following the steps in section A.1. So finding a  $\delta$  covering for  $\mathcal{F}_n$  is equivalent to finding an  $l_1 \delta$  covering for the space of coefficients in the tangent space using  $L_\infty$  norm. Let us have a closer look at the space of coefficients. We have  $\|v\| < \pi/4$  for tangent space representation of  $\Gamma$ , which is equivalent to  $\|c\|_2 \leq l_3$ , say. Therefore  $\mathcal{F}_n \equiv \{c \in \mathbb{R}^{k_n} : \|c\|_2 \leq l_3\} = \mathcal{C}$ , say. Then  $\mathcal{C} \subset \{c \in \mathbb{R}^{k_n} : \|c\|_\infty \leq l_4\} \equiv \{c \in \mathbb{R}^{k_n} : |c_i| \leq l_4 \forall i = 1, \dots, k_n\} = \mathcal{C}_0$ , say. Now  $\mathcal{C}_0$  is a compact set with  $\mathcal{C}$  as a compact subset. So the  $\delta$ -cover for  $\mathcal{C}$  would be a proper subset of the  $\delta$ -cover for  $\mathcal{C}_0$ . Therefore the covering number  $N$  for  $\mathcal{C}$  would be less than the covering number for  $\mathcal{C}_0$ . Since  $\mathcal{C}_0 \equiv \{[-l_4, l_4]^{k_n}\}$ , we have the covering number for  $\mathcal{C}_0$  as  $(\frac{2l_4}{l_1 \delta})^{k_n}$ . We obtain this by partitioning the interval  $[-l_4, l_4]$  into pieces of length  $l_1 \delta$  for each coordinate so that the partition of  $\mathcal{C}_0$  is reached through cross product. Then in each equivalent class of the partition of  $\mathcal{C}_0$  we will have  $\|c_1 - c_2\|_\infty \leq l_1 \delta$  which is equivalent to  $H(f_1, f_2) \leq \delta$ . So we have the metric entropy for  $\mathcal{F}_n = H(\cdot, \mathcal{F}_n) = H(u, \mathcal{F}_n) < k_n \log l/u$ , where  $l = 2l_4$  and  $u = l_1 \delta$ .

Now,

$$\int_{\epsilon^2/2^8}^{\sqrt{2}\epsilon} H^{1/2}\left(\frac{u}{l_3}, \mathcal{F}_n\right) du \leq \sqrt{k_n} \int \sqrt{\log(l_0/u)} du \leq \sqrt{k_n \log(M/\epsilon^2)} (\sqrt{2}\epsilon - \epsilon^2/256)$$

where  $l_0 = l_3 l$  and  $M = 2^8 l_0$ . For the existence of an  $\epsilon_n$  that satisfies Theorem 2 we need an  $\epsilon_n$  less than 1

that satisfies

$$\sqrt{k_n \log(M/\epsilon^2)}(\sqrt{2}\epsilon - \epsilon^2/256) \leq C_4 n^{1/2} \epsilon^2 \quad (\text{A.5})$$

But this inequality holds at 1— and hence there exists a smallest  $\epsilon_n < 1$  that satisfies A.5.  $\square$

### A.3 Proof of Theorem 1

Let  $C_1, \dots, C_4$ , be as in Theorem 2 and  $\epsilon_n$  be the smallest value of  $\epsilon$  satisfying condition A.4.  $\mathcal{F}_n$  be a sequence of approximating spaces,  $q$  be the corresponding  $\eta_n$ -sieve MLE. Let  $\delta_n$  be as defined in section A.

Define, for any  $0+ \leq \alpha \leq 1$ ;

$$\epsilon_n^*(\alpha) = \begin{cases} \epsilon_n, & \text{if } \delta_n(\alpha) < \frac{1}{4}C_1\epsilon_n^2, \\ (4\delta_n(\alpha)/C_1)^{1/2}, & \text{otherwise} \end{cases}$$

It follows from Wong and Shen [1995], that for any  $\alpha \in (0, 1]$ , if  $\delta_n(\alpha) < 1/\alpha$  and  $\eta_n \equiv 0$ , then

$$P(\|q^{1/2} - p_0^{1/2}\|_2 \geq \epsilon_n^*(\alpha)) \leq 5\exp(-C_2 n(\epsilon_n^*(\alpha))^2) + \exp(-\frac{1}{4}n\alpha C_1(\epsilon_n^*(\alpha))^2). \quad (\text{A.6})$$

Consider  $\alpha = 1$  in (A.6).  $\delta_n(1) = \inf_{q \in \mathcal{F}_n} \rho_1(p_0, q) = \inf_{q \in \mathcal{F}_n} \int p_0 g_1(p_0/q) \cdot = \inf_{q \in \mathcal{F}_n} \int p_0 \frac{(p_0 - q)}{q}$ .

Let  $P_0$  and  $F_p$  be the cdfs corresponding to the true density and the initial parametric estimate respectively.

Then we have  $\gamma = F_p^{-1} \circ P_0$  has the tangent space representation  $v_0$  obtained via exponential map of

$\sqrt{\gamma}$  satisfying  $\|v_0\| < \pi/4$ . This forces  $\cos(\|v\|) + \frac{\sin(\|v\|)}{\|v\|}v$  to be always positive. Now we look at

$\|p_0 - q\|_\infty$ , where  $q$  is the final warped maximum likelihood estimate of  $p_0$ . Let  $c_1$  be the coefficient vector

corresponding to  $p_0$ . Let  $c_2$  be the coefficient vector corresponding to  $q$ . Then  $c_1$  is infinite dimensional and

$c_2$  is just the first  $k_n$  coordinates of  $c_1$ , where  $k_n$  is the dimension of the coefficient vectors corresponding

to  $\mathcal{F}_n$ . We replace  $c_2$  by its corresponding infinite dimensional vector, i.e the infinite dimensional vector

with first  $k_n$  coordinates  $c_2$  and the rest coordinates all 0. Then we have  $\|p_0 - q\|_\infty < \|c_1 - c_2\|_\infty <$

$M\|c_1^{>k_n}\|_\infty \sim n^{\frac{-\beta}{2\beta+1}}$ , where  $\beta$  is the smoothness degree of  $p_0$ . Thus, assuming that  $k_n$  is large enough,

we have  $\|v_1\| < \pi/4$  corresponding to  $q$ . That implies  $\cos(\|v_1\|) + \frac{\sin(\|v_1\|)}{\|v_1\|}v_1 > 0$ , i.e.

$$\dot{\gamma}(t) = \left( \cos(\|v_1\|) + \frac{\sin(\|v_1\|)}{\|v_1\|}v_1 \right)^2(t) > 0 \quad \forall t \in [0, 1]$$

Also  $\dot{\gamma}(t)$  is continuous in  $t$  on a closed and bounded interval. So it attains its minima at some point  $t_0$

such that  $\dot{\gamma}(t) \geq \dot{\gamma}(t_0) > 0$  for all  $t \in [0, 1]$ . Thus it follows that  $q(t) > M_1 \dot{\gamma}(t_0) = d$ , say. Then we

have,  $\delta_n(1) = \int_0^1 \{p_0(p_0 - q)/q\}(t)dt < \|p_0 - q\|_\infty/d < M\|c_1^{>k_n}\|_\infty \sim n^{\frac{-\beta}{2\beta+1}}$ .

We have from equation A.5

$$\sqrt{k_n \log(M/\epsilon^2)}(\sqrt{2}\epsilon - \epsilon^2/256) \geq \sqrt{k_n \log(M/\epsilon^2)}c_0\epsilon \text{ where } c_0 = (\sqrt{2} - 1/256)$$

So for the smallest root we can solve the equation  $\sqrt{k_n \log(M/\epsilon^2)}c_0\epsilon = C_4n^{1/2}\epsilon^2$ . Let  $\epsilon_n$  be of the form  $\sqrt{M}n^{-\gamma}(\log n)^t$ ,  $\gamma > 0$ , and, let  $k_n = n^\Delta$ ,  $\Delta < 1$  Then

$$\log \frac{M}{\epsilon_n^2} = 2\gamma \log n + 2t \log \log n \geq (2\gamma + l_7) \log n, \text{ where } l_7 > 0$$

Therefore equating,  $n^{\Delta/2}\sqrt{2\gamma + l_7}\sqrt{\log n}\sqrt{M}n^{-\gamma}(\log n)^t$  with  $C_4Mn^{1/2}n^{-2\gamma}(\log n)^{2t}$ , we get  $\gamma = \frac{1}{2}(1 - \Delta)$ , and  $t = 1/2$ . Thus we have  $\epsilon_n = \sqrt{M}n^{\frac{-(1-\Delta)}{2}}\sqrt{\log n}$ . We take  $\Delta$  to be  $\frac{1}{2\beta+1}$  where  $\beta$  is the smoothness of the true density. Therefore  $\epsilon_n = \sqrt{M}n^{\frac{\beta}{2\beta+1}}\sqrt{\log n}$ . is the smallest value that satisfies the condition for Theorem 2. Therefore, using the definition given in Theorem 1, and using  $\alpha = 1$ , we get

$$\epsilon_n^* = \begin{cases} Mn^{-\beta/(2\beta+1)}\sqrt{\log n}, & \text{if } \delta_n(1) < \frac{1}{4}C_1M^2n^{-2\beta/(2\beta+1)}\log n, \\ (4\delta_n(1)/C_1)^{1/2}, & \text{otherwise} \end{cases}$$

## REFERENCES

- Ian S Abramson. On bandwidth variation in kernel estimates-a square root law. *The annals of Statistics*, pages 1217–1223, 1982.
- Ethan Anderes and Marc Coram. A general spline representation for nonparametric and semiparametric density estimates using diffeomorphisms. *arXiv preprint arXiv:1205.5314*, 2012.
- Ethan Anderes and Marc A Coram. Two-dimensional density estimation using smooth invertible transformations. *Journal of Statistical Planning and Inference*, 141(3):1183–1193, 2011.
- A Bhattacharya, D Pati, and DB Dunson. Latent factor density regression models. *Biometrika*, 97(1):1–7, 2010.
- Yeonseung Chung and David B. Dunson. Nonparametric bayes conditional distribution modeling with variable selection. *Journal of the American Statistical Association*, 104(488):1646–1660, 2009. URL <http://pubs.amstat.org/doi/abs/10.1198/jasa.2009.tm08302>.
- Miguel De Carvalho. Confidence intervals for the minimum of a function using extreme value statistics. *International Journal of Mathematical Modelling and Numerical Optimisation*, 2(3):288–296, 2011.

- Maria De Iorio, Peter Muller, Gary L. Rosner, and Steven N. MacEachern. An anova model for dependent random measures. *Journal of the American Statistical Association*, 99(465):205–215, 2004. URL <http://www.jstor.org/stable/27590366>.
- Hassan Doosti and Peter Hall. Making a non-parametric density estimator more attractive, and more accurate, by data perturbation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(2):445–462, 2016.
- David B Dunson and Ju-Hyun Park. Kernel stick-breaking processes. *Biometrika*, 95(2):307–323, 2008.
- David B. Dunson, Natesh Pillai, and Ju-Hyun Park. Bayesian density regression. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 69(2):pp. 163–183, 2007. ISSN 13697412. URL <http://www.jstor.org/stable/4623261>.
- Sam Efromovich. Orthogonal series density estimation. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(4):467–476, 2010.
- Michael D Escobar and Mike West. Bayesian density estimation and inference using mixtures. *Journal of the american statistical association*, 90(430):577–588, 1995.
- J. E Griffin and M. F. J Steel. Order-based dependent dirichlet processes. *Journal of the American Statistical Association*, 101(473):179–194, 2006. URL <http://pubs.amstat.org/doi/abs/10.1198/016214505000000727>.
- Peter Hall, Simon J Sheather, MC Jones, and James Stephen Marron. On optimal data-based bandwidth selection in kernel density estimation. *Biometrika*, 78(2):263–269, 1991.
- Bruce E Hansen. Nonparametric conditional density estimation. *Unpublished manuscript*, 2004.
- Nils Lid Hjort and Ingrid K Glad. Nonparametric density estimation with a parametric start. *The Annals of Statistics*, pages 882–904, 1995.
- Sonia Jain and Radford M Neal. A split-merge markov chain monte carlo procedure for the dirichlet process mixture model. *Journal of Computational and Graphical Statistics*, 2012.

- Maria Kalli, Jim E Griffin, and Stephen G Walker. Slice sampling mixture models. *Statistics and computing*, 21(1):93–105, 2011.
- Suprateek Kundu and David B Dunson. Latent factor models for density estimation. *Biometrika*, 101(3):641–654, 2014.
- Serge Lang. *Fundamentals of differential geometry*, volume 191. Springer Science & Business Media, 2012.
- Peter J Lenk. The logistic normal distribution for bayesian, nonparametric, predictive densities. *Journal of the American Statistical Association*, 83(402):509–516, 1988.
- Peter J Lenk. Towards a practicable bayesian nonparametric density estimator. *Biometrika*, 78(3):531–543, 1991.
- Tom Leonard. Density estimation, stochastic processes and prior information. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 113–146, 1978.
- Oleg Lepski et al. Adaptive estimation over anisotropic functional classes via oracle approach. *The Annals of Statistics*, 43(3):1178–1242, 2015.
- Qi Li and Jeffrey Scott Racine. *Nonparametric econometrics: theory and practice*. Princeton University Press, 2007.
- Matthew P Longnecker, Mark A Klebanoff, Haibo Zhou, and John W Brock. Association between maternal serum concentration of the ddt metabolite dde and preterm and small-for-gestational-age babies at birth. *The Lancet*, 358(9276):110–114, 2001.
- Steven N. MacEachern. Dependent nonparametric processes. *ASA Proceedings of the Section on Bayesian Statistical Science*, 1999. URL <http://aima.eecs.berkeley.edu/~russell/classes/cs294/f05/papers/maceachern-1999.pdf>.
- Steven N MacEachern and Peter Müller. Estimating mixture of dirichlet process models. *Journal of Computational and Graphical Statistics*, 7(2):223–238, 1998.
- J Steve Marron and Matt P Wand. Exact mean integrated squared error. *The Annals of Statistics*, pages 712–736, 1992.



- Peter Müller, Alaattin Erkanli, and MIKE West. Bayesian curve fitting using multivariate normal mixtures. *Biometrika*, 83(1):67–79, 1996.
- Andriy Norets and Justinas Pelenis. Bayesian modeling of joint and conditional distributions. *Journal of Econometrics*, 168:332–346, 2012.
- Murray Rosenblatt et al. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*, 27(3):832–837, 1956.
- S Saoudi, A Hillion, and F Ghorbel. Non-parametric probability density function estimation on a bounded support: Applications to shape classification and speech coding. *Applied Stochastic models and data analysis*, 10(3):215–231, 1994.
- S Saoudi, F Ghorbel, and A Hillion. Some statistical properties of the kernel-diffeomorphism estimator. *Applied stochastic models and data analysis*, 13(1):39–58, 1997.
- Simon J Sheather and Michael C Jones. A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 683–690, 1991.
- Anuj Srivastava and Eric P Klassen. *Functional and shape data analysis*. Springer, 2016.
- EG Tabak and Cristina V Turner. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.
- Esteban G Tabak and Giulio Trigila. Data-driven optimal transport. *Commun. Pure. Appl. Math.* doi, 10:1002, 2014.
- George R Terrell and David W Scott. Variable kernel density estimation. *The Annals of Statistics*, pages 1236–1265, 1992.
- Surya T Tokdar. Towards a faster implementation of density estimation with logistic gaussian process priors. *Journal of Computational and Graphical Statistics*, 2012.
- Surya T Tokdar, Yu M Zhu, Jayanta K Ghosh, et al. Bayesian density regression with logistic gaussian process and subspace projection. *Bayesian analysis*, 5(2):319–344, 2010.

- Bradley C Turnbull and Sujit K Ghosh. Unimodal density estimation using bernstein polynomials. *Computational Statistics & Data Analysis*, 72:13–29, 2014.
- Philippe Van Kerm et al. Adaptive kernel density estimation. *Stata Journal*, 3(2):148–156, 2003.
- V Vermehren and HM de Oliveira. Close expressions for meyer wavelet and scale function. *arXiv preprint arXiv:1502.00161*, 2015.
- Wing Hung Wong and Xiaotong Shen. Probability inequalities for likelihood ratios and convergence rates of sieve mles. *The Annals of Statistics*, pages 339–362, 1995.